Ethic Theory Moral Prac (2015) 18:565–575
DOI 10.1007/s10677-014-9540-x

When the Shape of a Life Matters

Stephen M. Campbell

Accepted: 23 September 2014 / Published online: 30 September 2014

Abstract It seems better to have a life that begins poorly and ends well than a life that begins well and ends poorly. One possible explanation is that the very shape of a life can be good or bad for us. If so, this raises a tough question: *when* can the shape of our lives be good or bad for us? In this essay, I present and critique an argument that the shape of a life is a *non-synchronic prudential value*—that is, something that can be good or bad for us in a way that is not good or bad for us at any particular time. After distinguishing two interpretations of 'the shape of a life', I argue that the first type of shape can be good or bad for us at particular moments while the other cannot be good or bad for us at all. This suggests that the shape of a life gives us no reason to posit non-synchronic prudential values.

Keywords Well-being · Prudential value · The good life · Shape of a life · Timing puzzle

There is a well-known puzzle about prudential value and time that arises from three sources. First, it seems plausible that something can be good or bad for us only if it can be good or bad for us at some particular time.[1] Second, there are certain sorts of things—including one's own death and some types of posthumous events—that seem as if they can be good or bad for an individual under some circumstances. Third, there appears to be no plausible story as to particular times when these things could be good or bad for an individual. For this reason, we may call them *hard cases*. Taken together, these three facts are puzzling. For we cannot coherently believe that nothing can be good or bad for us unless it can be so at one or more moments, that the hard cases can be good or bad for us, and that there are no moments at which they can be good or bad for us. At least one of these *prima facie* plausible beliefs must be rejected, but which one? This is the *Prudential Timing Puzzle*.[2]

[1] To achieve more precision, we must articulate distinct principles. One principle states that something is *intrinsically* good/bad for us only if it is intrinsically good/bad for us at some particular time. Another states that something is *instrumentally* good/bad for us only if it leads to or promotes something that is intrinsically good/bad for us at some particular time. (There will be further principles applying to other forms of derivative prudential value.)

[2] I borrow the phrase 'timing puzzle' from Steven Luper (2009), ch. 6, who discusses this puzzle in relation to death and posthumous events.

S. M. Campbell (✉)
Department of Medical Ethics and Health Policy, University of Pennsylvania, 3401 Market Street, Suite 320, Philadelphia, PA 19104, USA
e-mail: s.campbell80@gmail.com

Solving the puzzle requires finding a reasonable way to avoid inconsistency with respect to each hard case. Corresponding to the three sources are three strategies. The *synchronic strategy* challenges the third source by identifying particular times at which a given hard case can be good or bad for us. The *denial strategy* responds to the second source by denying that the hard case in question can have prudential significance after all. The *non-synchronic strategy* resists the first source by claiming that a hard case can be good or bad for us in a way that does not constitute a benefit or harm to us at any particular time.

One question worth asking is whether the non-synchronic strategy is a legitimate way of solving the Prudential Timing Puzzle. That strategy is unlike the other two in an important respect. Since virtually everyone believes that many things are good or bad for us at particular times (think of happiness and physical pain) and that many things are not good or bad for us at all (think of minute physical events that occur in distant parts of our universe), the synchronic and denial strategies do not require any alteration in our view of what sorts of things exist. The non-synchronic strategy is different. It posits *non-synchronic prudential values*: things that can be good or bad for us in a way that is not good or bad for us at any particular time.[3] In recent years, some philosophers have expressed an openness to the existence of such values,[4] but many of us come to the Prudential Timing Puzzle without any pre-theoretical conviction that non-synchronic values exist. And since the very idea of a non-synchronic value is peculiar, it is not unreasonable to wonder whether it is even intelligible that something could be good or bad for us non-synchronically.

My goal in this essay is to explore whether the so-called 'shape of a life' sheds new light on the existence of non-synchronic values and the viability of the non-synchronic strategy. The relationship between the shape of a life and the Prudential Timing Puzzle has not received much attention in the philosophical literature, but the connection is there. Like death and posthumous events, the shape of a life is a hard case. And it turns out to be an especially interesting hard case since a popular thought experiment about the shape of a life can be deployed to argue for the existence of non-synchronic values. It goes like this.[5] Imagine two ways that a person's life might unfold. It might go poorly for her during childhood but get progressively better over time. Alternatively, this individual might fare quite well early in life only to have things get progressively worse for her over time. Suppose that the two lives are of equal length and that the uphill life would begin as poorly as the downhill life ends and also end as wonderfully as the downhill life begins. Contemplating this pair of lives, many people are convinced that it is better for the person to have the uphill life than the downhill life. This has been dubbed the *Shape-of-a-Life Phenomenon*.[6] There are different ways that we might explain this normative truth (assuming that it is a truth), but one possibility is that the shape of a life itself has intrinsic prudential significance and that this explains why the uphill life is prudentially preferable to the downhill one.[7] Yet, since

[3] This definition leaves open the possibility that a single thing might be a non-synchronic value and a synchronic value. This can occur if something is non-synchronically good/bad for someone in one respect and synchronically good/bad for her in another.

[4] This tends to happen in one of two ways. It is sometimes said that a thing might be good or bad for us 'timelessly' or 'atemporally'. See, for instance, Luper (2009), 139; Bigelow et al. (1990), 121; Broome (2004), 47, 237–38; Bradley (2009), 74–78; Johansson (2013), 266–70; and Bramble (2014). Other times, the suggestion is that a thing might impact one's 'lifetime well-being' or 'diachronic well-being' without impacting that individual's 'synchronic well-being' at any time. (These terms are defined in Section 1.) See Bigelow et al. (1990); Velleman (1993); Glasgow (2013), 666; and Bramble (2014).

[5] What follows is adapted from an influential passage in Velleman (1993), 331. I quote the passage in Section 1.

[6] Feldman (2004), ch. 6; Portmore (2007), 21–24.

[7] This line of explanation has been entertained or endorsed by several philosophers, including Sen (1979), 470–71; Bigelow et al. (1990), 121–23, 137; Kamm (2003), 222–23; Temkin (2012), 111–12; and Glasgow (2013). Competing explanations of the Shape-of-a-Life Phenomenon are discussed in Slote (1983); Velleman (1993); Kamm (2003), 222–23; Feldman (2004), ch. 6; Portmore (2007), 21–24; Glasgow (2013), 669–80; and Dorsey (2014).

it is difficult to identify particular moments *when* the shape of our lives can be good or bad for us, it is tempting to conclude that the shape of a life must be a non-synchronic value.

These considerations point in the direction of a relatively straightforward, two-premise argument for the existence of non-synchronic values. I will call it the *Shape of a Life Argument*.

P1. The shape of a life can be intrinsically good or bad for a person.
P2. The shape of a life cannot be intrinsically good or bad for a person at any particular time.
C1. Thus, the shape of a life is a non-synchronic prudential value.
C2. Thus, non-synchronic prudential values exist.

Arguably, P1 is supported by our intuition about the above pair of lives as well as our intuitions about other kinds of shapes, and P2 is supported by our inability to identify moments at which this global feature of a life is good for us. If both premises are true, it follows that non-synchronic values exist.

This essay presents a critique of the Shape of a Life Argument. In §1, I distinguish two versions of the argument, which involve two distinct hard cases. On one version, the putative non-synchronic value is the shape of the temporal distribution of prudential impacts of first-order prudential values. In §2, I critique this argument by arguing that there are particular times at which this shape of one's life can matter prudentially. This is to challenge P2 and recommend the synchronic strategy for handling this hard case. On the other version, the putative non-synchronic value is the shape of the distribution of one's overall synchronic well-being levels over time. In §3, I challenge this version of the argument by raising doubts about the intelligibility of non-synchronic values. This is to cast doubt on P1 and highlight the denial strategy as a promising option. This critique does not settle the issue of whether non-synchronic values exist. But, if successful, it does serve to undermine two motivations for positing them.

1

What exactly is meant by 'the shape of a life'? This vague phrase admits of multiple interpretations. For the purpose of drawing out the two interpretations that will be the focus of this essay, let us examine David Velleman's influential presentation of the uphill-downhill life thought experiment[8]:

> Consider two different lives that you might live. One life begins in the depths but takes an upward trend: a childhood of deprivation, a troubled youth, struggles and setbacks in early adulthood, followed finally by success and satisfaction in middle age and a peaceful retirement. Another life begins at the heights but slides downhill: a blissful childhood and youth, precocious triumphs and rewards in early adulthood, followed by a midlife strewn with disasters that lead to misery in old age. Surely, we can imagine two such lives as containing equal sums of momentary well-being. Your retirement is as blessed in one life as your childhood is in the other; your nonage is as blighted in one life as your dotage is in the other. Yet even if we were to map each moment in one's life onto a moment of equal well-being in the other, we would not have shown these lives to be equally good. For after the

[8] It should be clarified that, while my discussion of the Shape of a Life Argument is inspired by Velleman's discussion, he does not present or discuss this argument nor does he talk much of 'shape'. Instead, he employs this thought experiment to argue against the view that 'well-being is additive', a view that implies that one's lifetime well-being necessarily equals the sum of her synchronic well-being levels. Velleman (1993), 331–32.

> tally of good times and bad times had been rung up, the fact would remain that one life gets progressively better while the other gets progressively worse; one is a story of improvement while the other is a story of deterioration. To most people, I think, the former story would seem like a better life-story—not, of course, in the sense that it makes for a better story in the telling or the hearing, but rather in the sense that it is the story of a [prudentially] better life.[9]

This passage sheds light on two ways of interpreting the phrase 'the shape of a life'. The first is implied by Velleman's talk of 'momentary well-being'. In the philosophical literature, it is now customary to distinguish three forms of well-being. *Synchronic* (*or momentary*) *well-being* is a comprehensive measure of how well one is doing or faring at a particular time. *Diachronic well-being* is a measure of how well one fares over some period of time. *Lifetime well-being* concerns how well one's life as a whole goes for her and is supposed to capture all of the ways in which things are good or bad for an individual from conception to death—and, if applicable, beyond. On one interpretation, the phrase 'the shape of a life' refers to the shape of an individual's synchronic well-being curve—that is, the distribution of his or her synchronic well-being levels over time. I will refer to this as a life's *synchronic well-being shape*.

Next, notice that Velleman's initial description of the two lives draws our attention to more basic prudential values and disvalues, such as happiness, misery, achievements, and setbacks. These are *first-order prudential values* since they do not supervene upon other prudential values in the way that synchronic well-being shape does. They are also *synchronic prudential values* since, presumably, they are good or bad for people at particular times. Let the phrase *first-order synchronic well-being* refer to how well an individual is faring at a particular time only with respect to first-order, synchronic prudential values. It is natural to think that having an upward trend in terms of first-order synchronic well-being is better for a person than having a downward trend. On a second interpretation, 'the shape of a life' refers to the shape of the distribution of an individual's first-order synchronic well-being levels over time. Call this a life's *first-order shape*.

It might be wondered whether it makes sense to distinguish these two shapes. On many normative views, synchronic well-being shape and first-order shape will always coincide. For instance, if hedonism is true, then synchronic well-being shape and first-order shape will always match the 'hedonic shape' of a life, which reflects how well one does over time in terms of pleasure and pain. Yet, the question of whether these two shapes ever diverge is a substantive normative matter. One can coherently believe that there are *second-order* prudential values (such as the first-order shape of a life) that are good or bad for an individual at particular times. If there are second-order, synchronic prudential values of this sort, a life's first-order shape and synchronic well-being shape will diverge in some cases.

The first-order shape of a life and the synchronic well-being shape of a life represent two hard cases. They yield two versions of the Shape of a Life Argument and two distinct ways of trying to establish the existence of non-synchronic values. As we shall see, the two arguments call for different critiques.

2

One version of the Shape of a Life Argument concerns the first-order shape of a life, or the shape of the distribution of one's first-order synchronic well-being levels over time.

[9] Ibid., 331.

P1. The first-order shape of a life can be intrinsically good or bad for a person.
P2. The first-order shape of a life cannot be intrinsically good or bad for a person at any particular time.
C1. Thus, the first-order shape of a life is a non-synchronic prudential value.
C2. Thus, non-synchronic prudential values exist.

For simplicity, suppose that the only first-order synchronic prudential values are happiness and knowledge. On this supposition, the first-order shape of an individual's life is the shape of the curve representing how well she is doing over time only with respect to happiness and knowledge. P1 asserts that the shape of this curve can have prudential significance. I have no real objection to this premise. All else being equal, it does strike me as being prudentially preferable to have a life that gets progressively better in terms of happiness and knowledge than a life that gets progressively worse in those respects.[10]

This leaves P2, which asserts that the first-order shape of a life cannot be good or bad for an individual at any particular time. The claim is that this putative second-order prudential value cannot have a positive or negative impact on an individual's level of synchronic well-being at any moment. If we endorse P1, should we accept this premise? One possible justification of P2 draws upon a suggestion found in David Velleman that 'how a person is faring at a particular moment is a temporally local matter'.[11] He elaborates:

> [M]omentary well-being is ordinarily conceived as a temporally local matter, determined by a person's current circumstances, whether experienced or unexperienced. We think of a person's current well-being as a fact intrinsic to the present, not as a relation that he currently bears to his future.[12]

In his discussion of the Shape-of-a-Life Phenomenon, Douglas Portmore characterizes this view of synchronic well-being as follows:

> Momentary well-being is the welfare value that some momentary segment of one's life would have if that segment existed alone, apart from any relationship it has with other segments of one's life.[13]

The suggestion is that one's synchronic well-being level at some moment is solely a function of what is happening at that very moment.[14] If synchronic well-being is temporally local in this sense, it follows that the first-order shape of a life—along with any other state of

[10] Glasgow (2013) defends something close to this view. However, his focus is on the shape of one's 'non-relational [synchronic] well-being', which is a measure of how well a person is doing only in virtue of states of affairs that obtain at that very time. First-order synchronic well-being, as I have defined it, need not screen off the prudential impacts of states of affairs obtaining at other times, though it does screen off the impacts of second-order prudential values.

[11] Velleman (1993), 331.

[12] Ibid., 339.

[13] Portmore (2007), 21. I should emphasize that Portmore offers this characterization in the process of explicating Velleman and does not seem to be committed to it as a claim about synchronic well-being. Later in the same essay (26), he acknowledges the possibility that 'welfare value is relational' in such a way that 'the value of our past sacrifices can be affected (or determined) by subsequent events'. Glasgow (2013), 666, offers a very similar characterization of (non-relational) 'momentary well-being', though he immediately acknowledges the possibility that 'some moments' value can be at least partly relational'.

[14] This view is closely related to the principle that Ben Bradley dubs 'Internalism', according to which 'The intrinsic value of a time for a person is determined entirely by the value atoms obtaining at that time.' Bradley provides a positive argument for Internalism at (2009), 19. For two critiques of that argument, see Johansson (2013), 262–63, and Dorsey (2013), 167–69. See also McMahan (2002), 180, and Portmore (2007), 25–26.

affairs that is not wholly confined to a moment—cannot affect synchronic well-being since it is not a temporally local property.

Before proceeding with my critique of this proposal, there is a potential ambiguity that deserves mention. In line with Velleman's suggestion, it is true that the term 'well-being' is sometimes used to refer to a temporally local dimension of a person's condition. For instance, one's 'mental well-being' and 'physical well-being' at a time is typically understood to be solely a function of one's mental or physical state at that very time.[15] This restrictive use of 'well-being' (which seems intimately related to our talk of 'health') is useful in some domains, such as medicine and public policy, where we have reason to monitor alterations in some dimension of people's situation over time. Perhaps one can extend this usage and talk of 'overall well-being' at a time, where this is understood to be a function only of the state of the world at that time, including those aspects of the world that determine mental well-being, physical well-being, spiritual well-being, and so on. If the term 'well-being' is used in this way, there is no reason to resist the claim that well-being is a temporally local phenomenon.

That said, our present interest is in a use of 'well-being' that presupposes a tight connection between well-being (how well S is doing) and prudential value (what is good or bad for S). On this usage, something contributes to your level of well-being if and only if it is *good for you*, or *benefits* you, and something detracts from your level of well-being if and only if it is *bad for you*. Something contributes to (or detracts from) your level of well-being *at a time* if and only if it is good (or bad) for you at that time. This sense of 'well-being' is widely invoked by philosophers and has a nice fit with everyday conversation. When you ask another person 'How are you today?' or 'How are you doing right now?', this can be an invitation to share news of anything that is good or bad for the person at that time. For the purpose of analyzing the Prudential Timing Puzzle, the question we must ask is whether synchronic well-being, in this latter sense, is 'a fact intrinsic to the present'.

With this clarification in mind, it seems clear that how well a person is doing at a particular time is *not* merely a temporally local matter. Ask someone how they are doing today, and you are quite likely to receive an answer that draws upon what happened last week or what will happen tomorrow. In some cases, people may feel unable to say how well they are doing in the present without having certain information about the past or future. To appreciate this fact, we need only consider certain cases in which the "meaning" of our activities depends on future events and, arguably, affects how well we are doing at the very time that we are engaged in them.[16] Imagine a person whose dream is to practice law and who is currently undergoing much stress and devoting substantial amounts of time, money, and energy to obtain a law degree. The meaning and significance of her present activity partly depends on future events. If she will ultimately attain her goal, her current efforts may be accurately described as laying the groundwork for her future career. However, if she will ultimately fail to secure the degree, and if there will be no real benefits reaped from this eventual failure or the process leading up to it, then her current efforts may be accurately characterized as wasted effort. It seems fair to say that, all else being equal, a person who is wasting her time, money, and efforts is worse off than a person who is engaged in similar activities without any such waste. In this way, the meaning

[15] Bigelow et al. (1990) argue against this view. They make the case that the properties that determine one's mental well-being and physical well-being 'will include a host of relational properties that link the person to events or states, some of which will be at that very same time, but many of which will be at other times' (133–34).

[16] I borrow this talk of 'meaning' from Velleman (1993), though he denies that meaning, at least when it is determined by past or future events, can influence synchronic well-being (339–40). On this matter, I side with McMahan (2002), 179–80, and Dorsey (2014), who contend that meaning can influence synchronic well-being. See also Portmore (2007), 25–26, who takes this latter view seriously without committing himself to it.

of our activities seems relevant to determining how well we are faring at particular times. In a similar vein, it is sometimes thought that our activities in life are meaningless or absurd because we as individuals will eventually die or because the human species will eventually go extinct.[17] Such views are often accompanied by the conviction that engaging in meaningless or absurd activities is bad for us at the very times we are engaging in them.

There are also examples in which the meaning and prudential status of activities seems to be determined partly by past facts. When Edmund Hillary and Tenzing Norgay reached the summit of Mount Everest in 1953, this was an achievement that was presumably very good for them at that time. But the magnitude of their achievement depended on facts about the past—particularly, the fact that no one before them had managed to accomplish this feat. Of course, there are other examples in which the fact that others *have* done what one is now doing can imbue that activity with greater prudential significance. The benefit of taking part in rituals, ceremonies, and pilgrimages often seems enhanced by the fact that so many others have engaged in these very same activities before us and, perhaps also, that so many others will do so after we are gone. It is partly in virtue of facts about other times that engagement in some activity can qualify as *participation in a tradition*. Thus, past facts can also affect the meaning of our activities and, in turn, affect their prudential significance.

Lastly, think of our identities. It is not uncommon for people to consider themselves fortunate or unfortunate in virtue of some aspect of who they are. One can feel pride about being raised as a Unitarian or being of Chinese descent. One might feel ashamed about being descended from a traitor, or of belonging to a society that is guilty of atrocities. But such identities, while they do seem to be good or bad for individuals at particular moments, are not temporally local phenomena. They depend on facts about the past—and, in some cases, events that occurred well before the person in question was even conceived.

These examples suggest that ordinary thinking about what is good or bad for us and how well or poorly we are doing at a given time does not screen off facts about other times. It is not important, for present purposes, that the reader endorses any of the above normative views. What matters is that the above judgments are intelligible and represent views that a person who has competence with our normative concepts might be led to endorse. That is sufficient to establish that synchronic well-being, in the sense that interests us, is not a temporally local phenomenon. This removes one obstacle from thinking that the first-order shape of a life is good or bad for us at particular times.

What remains is to identify particular moments when the first-order shape of a life might have prudential significance. For present purposes, I will focus on the downhill first-order shape that is described in Velleman's thought experiment. If it is bad for a person that her first-order synchronic well-being levels get progressively worse from birth to death, I submit that this is bad for her at *every* moment of her life. If it were known during her childhood that her level of happiness, achievement, etc. would steadily decline over time, it would be sensible to consider her, at that very time, unfortunate on account of this fact. It would also be sensible to judge this bad for her at moments in the middle of her life inasmuch as things have been going downhill and will only continue to do so. The same may be said of moments near the end of her life given that she has had a life of constantly diminishing first-order synchronic well-being.

It is worth emphasizing that these judgments concern what is bad for a person in a fairly limited respect. Presumably, a person living that sort of downhill life is not *all-things-considered* badly off at moments during her happy, achievement-filled childhood, even if

[17] A variation on this view is explored in Scheffler (2013), who observes that we would lose motivation to engage in many of our current projects and activities if we came to believe that the human race will be wiped out in the near future.

she is poorly off in the particular respect that her life will get progressively worse in terms of these prudential values. And even if she will be all-things-considered unfortunate at moments near the end of her life, this will be primarily due to her low level of first-order prudential goods. If that is correct, the prudential significance of the first-order shape of a life at a particular moment is relatively minor in comparison with the significance of the first-order, prudential values obtaining at that time.[18] This might help to explain why people have been led to think that the first-order shape of a life has no prudential significance at particular moments.

I do not mean to insinuate that, in every case, the first-order shape of a life will have a uniform impact at each moment of a person's life. Imagine a life in which one's first-order synchronic well-being deteriorates at an increasing rate over time. Plausibly, this first-order shape is worst for the person at moments near the end of her life where the downward slope is steepest. Nor should we assume that the shape of a life is necessarily good or bad for an individual at every moment of her life. In some cases, a portion of the first-order shape of one's life might be prudentially neutral, while another portion might be significantly good or bad for her.

A complete account of the prudential significance of the first-order shape of a life would tell us, for every possible first-order shape of a life, how good or bad it is for the individual in question at each moment of time. Given the vast array of first-order shapes that lives may have, as well as the practical difficulty of specifying degrees of prudential value, providing a complete account is likely to prove a practical impossibility. Thankfully, there is little or no practical benefit to be derived from formulating such an account. Our interest in the momentary prudential significance of the first-order shape of a life is grounded in the desire to avoid inconsistency. In most contexts, we have little practical use for detailed information about precisely when and to what degree the shape of one's life is good for her. I suspect that this point generalizes to all of the hard cases. We have no practical need for a complete account of precisely when and to what extent death, posthumous events, etc. are good or bad for us. We just need to avoid contradictions.

I have only considered a few first-order shapes that a life might have. Even so, this limited examination suggests that there are moments when the first-order shape of a life might be good or bad for a person. This means that benefits and harms that come to us from the first-order shape of our lives can be captured by our synchronic well-being levels. Intuitively, this is far more plausible than the view that there is no particular time when a person can be benefited or harmed by the first-order shape of her life. We therefore have reason to reject P2 in the first version of the Shape of a Life Argument.

3

This brings us to the second version of the Shape of a Life Argument, which pertains to synchronic well-being shape, or the shape of one's synchronic well-being curve.

P1. The synchronic well-being shape of a life can be intrinsically good or bad for a person.
P2. The synchronic well-being shape of a life cannot be intrinsically good or bad for a person at any particular time.
C1. Thus, the synchronic well-being shape of a life is a non-synchronic prudential value.
C2. Thus, non-synchronic prudential values exist.

[18] Glasgow (2013), 681, defends a similar point.

On this version of the argument, it is difficult to see how one might sensibly reject P2. If P1 is false, then P2 is certainly true. For if the synchronic well-being shape of a life cannot be good or bad for a person, then it cannot be good or bad for a person at any particular time. If P1 is true, then P2 is very plausibly true. There are various absurdities that arise if having a specific synchronic well-being shape could be good or bad for us at particular times. Consider the case of Zack, who is convinced that Charlie Chaplin led an ideal life and that it is good for a person to have a life resembling Chaplin's in various ways. For instance, Zack believes that it substantially benefits a person to have a synchronic well-being curve exactly like Chaplin's.[19] On the face of it, this view seems intelligible. For the sake of argument, suppose it is true. Could having Chaplin's exact synchronic well-being curve be substantially good for a person *at a particular time*? It seems not. To see why, think of Chaplin himself. Suppose that having his exact synchronic well-being curve benefited him at some particular time. Was that benefit already factored into the curve? If so, this implies that he was mysteriously bootstrapped into having his exact synchronic well-being curve with the aid of the benefit of having that exact synchronic well-being curve! This is a very bizarre result. If the benefit was not factored in, a contradiction emerges. For, then, Chaplin had an exact synchronic well-being curve that resulted in a benefit that showed up somewhere in his synchronic well-being curve, with the result that did not have that exact curve. If P2 is true, these absurdities are avoided.

Embracing P2 leaves two options. On the one hand, we might accept P1 and maintain that non-synchronic values exist. This is to embrace the non-synchronic strategy with respect to the synchronic well-being shape of a life. On the other hand, we might reject P1 and deny that the synchronic well-being shape of a life can be good or bad for a person after all. This is to adopt the denial strategy for this hard case.

One way of motivating P1 appeals to those of us who are inclined to see the first-order shape of a life as having intrinsic prudential significance.[20] Why, it may be asked, should we care only about the shape of our *first-order* synchronic well-being levels? That seems arbitrary. If there are second-order synchronic prudential values, it stands to reason that we should also care about the shape that results from their impact on our lives. 'Second-order' does not mean second-class. To avoid having concerns that arbitrarily screen off some synchronic prudential values, we should focus on the shape of our *overall* synchronic well-being levels (which capture impacts by *any* synchronic prudential values) rather than some restricted form of synchronic well-being.

While this argument appears to have some force, it cannot succeed if its conclusion is *unintelligible*. Might P1 be unintelligible? One might argue that it is not by appealing to P1's *prima facie* intelligibility. It might be said: 'Since Zack's belief (that having Chaplin's exact synchronic well-being curve is a substantial benefit) appears to be an intelligible substantive view, we should assume that it is an intelligible view—at least until proven otherwise'.

This reasoning seems specious. Suppose that Zack also thinks that having Chaplin's exact lifetime well-being level substantially benefits a person. On the face of it, this also seems like an intelligible view. Supposing that it is true, we find ourselves in familiar territory. Is this benefit already factored into Chaplin's lifetime well-being level or not? If so, he was mysteriously bootstrapped into having his exact lifetime well-being level. If not, we face a contradiction. These are unpalatable results. Similar to the situation above, two options remain. We

[19] To be clear, Zack does not think that having this synchronic well-being curve is good for a person *because* Chaplin had that curve. Rather, he considers having that synchronic well-being shape to be beneficial, in and of itself. Thus, on his view, having that synchronic well-being shape was intrinsically good for Chaplin himself.

[20] The argument that follows can be adapted for other forms of restricted well-being, such as non-relational synchronic well-being (see note 10). But I will continue to focus my discussion on first-order synchronic well-being.

might conclude that some things are good or bad for a person without affecting his or her lifetime well-being level, or we might deny that Zack's view about having Chaplin's exact lifetime well-being level is an intelligible substantive view after all. However, a moment's reflection reveals that the first option is not a viable one in this case. Lifetime well-being is understood to be a measure that captures *all* of the ways in which things go well or poorly for us; it is a comprehensive measure of everything that is prudentially significant in a life. Therefore, we cannot take the first option, for there is no dimension of well-being that is not captured by lifetime well-being. Nor should we allow for bootstrapping or contradiction. We should embrace the second option and simply deny that Zack's view is intelligible after all. This shows that there are unintelligible views about well-being that can appear intelligible at first glance. Zack's view about the benefit of having Chaplin's exact lifetime well-being level is one of them, and it is perfectly possible that Zack's similar view about the benefit of having Chaplin's exact synchronic well-being curve is another. The mere *prima facie* intelligibility of the latter view is not sufficient to justify the assumption of intelligibility.

There is good reason to doubt the intelligibility of P1. If synchronic well-being shape is something that can be intrinsically good or bad for a person, then, for reasons already discussed, it is plausibly a non-synchronic prudential value. But the very idea of non-synchronic benefits and harms is quite strange.[21] To draw out the oddness, imagine this. You receive a note from a trusted source. It reads: 'I've just learned of something that is very, very bad for you. A serious non-synchronic harm. More details later.' Initially, you are distraught by this news. Yet, as you reflect upon the nature of non-synchronic evils, you take some comfort in the fact that this thing is not bad for you in the least *at the present moment*. Nor, you next realize, was it bad for you at all at the previous moment, or the moment before that, or the moment before that… Nor will it be bad for you in the next moment, or the moment after that, or the moment after that… Indeed, you come to recognize, it is not bad for you *at any moment in all of history*! It is as if you move through time with complete immunity from the badness of this thing. You find this to be a substantial consolation and, upon further reflection, decide that you do not mind being harmed non-synchronically.[22]

This little tale is useful for highlighting the strangeness of the idea of non-synchronic prudential values. Since non-synchronic harms would not be bad for us at any particular moment, it does not seem unreasonable for a person to react as you do in the imagined scenario. Indeed, it is not obviously unreasonable that one should be altogether indifferent to non-synchronic harms and always prefer to "suffer" a serious non-synchronic harm rather than suffer a minor synchronic harm. But all of this flies in the face of our ordinary understanding of what harm is. Serious harm is not the sort of thing that can be reasonably ignored by the one who undergoes it. The fact that it seems an open question whether serious non-synchronic harms can be ignored and whether they can be preferred over minor synchronic harms casts serious doubt on the intelligibility of that idea. This, in turn, gives us reason to doubt the intelligibility of P1.

Rejecting P1 in the second version of the Shape of a Life Argument does not seem very costly. While it is true that many people have the intuition that an uphill life is prudentially preferable to a downhill one, it is unlikely that this intuition concerns the synchronic well-

[21] Here, I am in full agreement with Jennifer Hawkins (2014), 535–36. She notes that the idea of a non-synchronic prudential value (which she calls a 'timeless good') 'is really incredibly odd when we start to think about it' and ultimately concludes that we should only posit such values as a last resort.

[22] Granted, news of a non-synchronic harm might be distressing if, and to the extent that, it provides evidence that something is synchronically bad for you. But my present concern is one's attitudes toward the non-synchronic harm itself.

being shape of a life. The idea of a synchronic well-being curve—that is, a comprehensive measure of *all* of the ways that things are good or bad for a person at particular moments of her life—is, after all, a rather abstract notion. Typically, when we think about a life getting progressively better or worse over time, we are focused on first-order synchronic values. We are thinking about a life improving or degrading in terms of one's happiness, life-satisfaction, social connections, knowledge, achievements, etc. In short, we are attuned to first-order shape. First-order shape and synchronic well-being shape are distinct phenomena, and it is perfectly coherent to think that the first-order shape of a life can be prudentially significant while the synchronic well-being shape of a life cannot. This, I propose, is what we should think.

4

When does the shape of a life matter, prudentially speaking? That depends on what is meant by 'shape of a life'. If it refers to the shape of the distribution of one's first-order synchronic well-being levels over time, I have suggested that this shape of a life can be intrinsically good or bad for us at particular moments in our lives. If it refers to the shape of the distribution of an individual's synchronic well-being levels over time, I have challenged the idea that this shape can be intrinsically good or bad for us at all. Either way, it appears that the shape of a life gives us no reason to believe in non-synchronic prudential values.

Acknowledgments For their helpful comments, I wish to thank Alex Sarch, Sven Nyholm, Jason Konek, Elizabeth Anderson, Allan Gibbard, Peter Railton, Sarah Buss, Anne Baril, David Wiens, Billy Dunaway, Ben Bramble, Dale Dorsey, Ralf Bader, two anonymous reviewers of ETMP, and the audience at the 2013 APA Pacific Division.

References

Bigelow J, Campbell J, Pargetter R (1990) Death and Well-Being. Pac Philos Q 71:119–40
Bradley B (2009) Well-being and death. Oxford University Press, Oxford
Bramble B (2014) Whole-life welfarism. Am Philos Q 51:63–74
Broome J (2004) Weighing lives. Oxford University Press, Oxford
Dorsey D (2013) Desire-satisfaction and welfare as temporal. Ethical Theory Moral Pract 16:151–71
Dorsey D (2014) The Significance of a Life's Shape. Ethics (forthcoming)
Feldman F (2004) Pleasure and the good life. Oxford University Press, Oxford
Glasgow J (2013) The shape of a life and the value of loss and gain. Philos Stud 162:665–82
Hawkins J (2014) Well-being, time, and dementia. Ethics 124:507–42
Johansson J (2013) The timing problem. In: Bradley B et al (eds) Oxford handbook of philosophy and death. Oxford University Press, Oxford, pp 255–73
Kamm FM (2003) Rescuing Ivan Ilych: how we live and how we die. Ethics 113:202–33
Luper S (2009) The philosophy of death. Cambridge University Press, Cambridge
McMahan J (2002) The Ethics of Killing: Problems at the Margins of Life. Oxford University Press, Oxford
Portmore D (2007) Welfare, achievement, and self-sacrifice. J Ethics Soc Philos 2:1–28
Scheffler S (2013) Death and the Afterlife. Oxford University Press, Oxford
Sen A (1979) Utilitarianism and welfarism. J Philos 76:463–89
Slote M (1983) Goods and virtues. Oxford University Press, Oxford
Temkin L (2012) Rethinking the good. Oxford University Press, Oxford
Velleman JD (1993) Well-being and time. In: Fischer JM (ed) The metaphysics of death. Stanford University Press, Stanford, pp 329–57